

Behind the Report—Testing Addendum:

Accelerate Compute-Intensive Financial Workloads

This document provides the system-configuration details and step-by-step procedures that Prowess Consulting used to perform benchmark testing on two Dell™ platforms running different generations of AMD EPYC™ processors:

- Dell™ PowerEdge™ R6525 server powered by an AMD EPYC 7702 processor
- Dell PowerEdge R6625 server powered by an AMD EPYC 9554 processor

For the full analysis, read the report, "[Accelerate Compute-Intensive Financial Workloads](#)."

Testing was concluded on February 13, 2023.

Server Configurations

	Dell™ PowerEdge™ R6525	Dell™ PowerEdge™ R6625
Processor	AMD EPYC™ 7702	AMD EPYC™ 9554
Number of CPUs	2	2
Cores per CPU	64	64
Total Cores/Threads	128/256	128/256
Frequency (Base/SCT/MCT)	2.0 GHz (max 3.9 GHz)	3.1 GHz (max 4.4 GHz)
Storage Controller 01	Dell™ PowerEdge™ RAID Controller (PERC) S150	Dell™ PowerEdge™ RAID Controller (PERC) H755
Disk	1.92 TB Intel® SSDSC2KB019T8R	1.92 TB Dell™ NVMe® v2 AGN RI U.2
Number of Disks	2	1
Installed Memory	512 GB	768 GB
Memory DIMM	Hynix® HMA84G7CJR4N-XN	Micron® MTC20F2085S1Rc48BA1
Memory Speed	3,200 MT/s	4,800 MT/s
Number of Memory DIMMs	16	24
BIOS Version	2.7.3	1.13
Software		
Operating System (OS)	Red Hat® Enterprise Linux®	
OS Version	8.7	
OS Kernel	4.18.0-425.10.1el8_7.x86_64	
Benchmarking Test Suite	AI Benchmark	

Test Procedures

Prowess Consulting used two systems for testing, one powered by 2nd Generation AMD EPYC processors and the other powered by 4th Generation AMD EPYC processors, to compare the systems' performance in AI Benchmark.

This testing was performed on hardware with Red Hat® Enterprise Linux® 8.7 installed.

Configure and Test the AI Benchmark Environment

The following steps outline configuring the benchmark environment and running the benchmark for testing.

1. Run the following commands to install Extra Packages for Enterprise Linux (EPEL) and Python® Pip:

```
subscription-manager repos --enable codeready-builder-for-rhel-8-$(arch)-rpms
dnf install https://dl.fedoraproject.org/pub/epel/epel-release-latest-8.noarch.rpm
yum repolist
yum install python3-pip
```

2. Run the following command to install Python 3.9.6:

```
wget https://www.python.org/ftp/python/3.9.6/Python-3.9.6.tgz
tar -xzvf Python-3.9.6.tgz
sudo yum -y install make
cd Python-3.9.6 && ./configure --enable-optimizations
make
make altinstall
```

3. Run the following command to install Anaconda® and Conda®:

```
sudo yum -y install anaconda
sudo yum -y install conda
```

4. Run the following command to initialize the Conda environment:

```
conda init bash
<exit><relog>
```

5. Run the following command to install Java and the Scala build tool (sbt):

```
sudo yum -y install java-1.8.0-openjdk java-1.8.0-openjdk devel
curl -L https://scala-sbt.org/sbt-rpm.repo > sbt-rpm.repo
mv sbt-rpm.repo /etc/yum.repos.d/
yum update
yum -y install sbt
```

6. Run the following command to install extra dependencies required for the benchmark:

```
yum install libglvnd-glx libsndfile libsndfile-devel
```

7. Download and extract AI Benchmark.

8. Navigate to the extracted AI Benchmark directory.

9. Run the following command to prepare the AI Benchmark workload environment:

```
chmod -R +x ./*
./setup-python.sh
```

10. Fix NumPy® and TensorFlow™ issues as the setup script installs the incorrect versions of NumPy and TensorFlow.

- a. Run the following commands to uninstall and install NumPy and TensorFlow to get the latest versions:

```
conda activate ../tpcx-ai-v1.0.2/lib/python-venv-ks
pip uninstall numpy
pip uninstall tensorflow
pip install tensorflow
pip install numpy
```

11. Tune the AI Benchmark environment:

- a. Run the following command to modify the setenv.sh file:

```
vi setenv.sh
```

- b. Update the **TPCxAI_SCALE_FACTOR** setting to **10**.
c. Update the **TPCxAI_SERVING_THROUGHPUT_STREAMS** setting to **100**.
d. Press the **Esc** key, and then run the following command to save the setenv.sh file:

```
:wq !
```

- e. Run the following command to set the number of **OpenMP Threads** to **128**:

```
export OMP_NUM_THREADS=128
```

12. Run the following command to validate AI Benchmark:

```
./TPCx-AI_validation.sh
```

13. In order to perform the benchmark testing, verify that the validation run returns zero errors, and then run the AI Benchmark test three times with the following command:

```
./TPCx-AI_Benchmark-run.sh
```



The analysis in this document was done by Prowess Consulting and commissioned by Dell Technologies.

Results have been simulated and are provided for informational purposes only. Any difference in system hardware or software design or configuration may affect actual performance.

Prowess Consulting and the Prowess logo are trademarks of Prowess Consulting, LLC..

Copyright © 2023 Prowess Consulting, LLC. All rights reserved.

Other trademarks are the property of their respective owners.